

Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach

By FRANK SMETS AND RAFAEL WOUTERS*

Using a Bayesian likelihood approach, we estimate a dynamic stochastic general equilibrium model for the US economy using seven macroeconomic time series. The model incorporates many types of real and nominal frictions and seven types of structural shocks. We show that this model is able to compete with Bayesian Vector Autoregression models in out-of-sample prediction. We investigate the relative empirical importance of the various frictions. Finally, using the estimated model, we address a number of key issues in business cycle analysis: What are the sources of business cycle fluctuations? Can the model explain the cross correlation between output and inflation? What are the effects of productivity on hours worked? What are the sources of the “Great Moderation”? (JEL D58, E23, E31, E32)

A new generation of small-scale monetary business cycle models with sticky prices and wages (the New Keynesian or New Neoclassical Synthesis (NNS) models) has become popular in monetary policy analysis.¹ Following Smets and Wouters (2003), this paper estimates an extended version of these models, largely based on Lawrence J. Christiano, Martin Eichenbaum, and Charles L. Evans (CEE, 2005), on US data covering the period 1966:1–2004:4, and using a Bayesian estimation methodology. The estimated model contains many

shocks and frictions. It features sticky nominal price and wage settings that allow for backward inflation indexation, habit formation in consumption and investment adjustment costs that create hump-shaped responses of aggregate demand, and variable capital utilization and fixed costs in production. The stochastic dynamics is driven by seven orthogonal structural shocks. In addition to total factor productivity shocks, the model includes two shocks that affect the intertemporal margin (risk premium shocks and investment-specific technology shocks), two shocks that affect the intratemporal margin (wage and price mark-up shocks), and two policy shocks (exogenous spending and monetary policy shocks). Compared to the model used in Smets and Wouters (2003), there are three main differences. First, the number of structural shocks is reduced to the number of seven observables used in estimation. For example, there is no time-varying inflation target, nor a separate labor supply shock. Second, the model features a deterministic growth rate driven by labor-augmenting technological progress, so that the data do not need to be detrended before estimation. Third, the Dixit-Stiglitz aggregator in the intermediate goods and labor market is replaced by the more general aggregator developed in Miles S. Kimball (1995). This aggregator implies that the demand elasticity of differentiated goods and labor depends on their relative price. As shown in Eichenbaum and Jonas Fischer (forthcoming), the introduc-

* Smets: European Central Bank/CEPR/University of Ghent, Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany (e-mail: frank.smets@ecb.int); Wouters: National Bank of Belgium/Université Catholique, Boulevard de Berlaimont 5, 1000 Brussels, Belgium (e-mail: rafael.wouters@nbb.be). The views expressed are solely our own and do not necessarily reflect those of the European Central Bank or the National Bank of Belgium. We thank seminar participants and discussants at the 2003 ECB/IMOP Workshop on Dynamic Macroeconomics, the Federal Reserve Board, Princeton University, the Federal Reserve Banks of St. Louis and Chicago, the 2004 ASSA meetings in San Diego, Köln University, Humboldt University, the European Central Bank, the Bank of Canada/Swiss National Bank/Federal Reserve Bank of Cleveland Joint Workshop on Dynamic Models Useful for Policy, and, in particular, Frank Schorfheide, Fabio Canova, Chris Sims, Mark Gertler, and two anonymous referees for very useful and stimulating comments.

¹ See Marvin Goodfriend and Robert G. King (1997), Julio J. Rotemberg and Michael Woodford (1995), Richard Clarida, Jordi Galí, and Mark Gertler (1999) and Woodford (2003).

tion of this real rigidity allows us to estimate a more reasonable degree of price and wage stickiness.

The objectives of the paper are threefold. First, as the NNS models have become the standard workhorse for monetary policy analysis, it is important to verify whether they can explain the main features of the US macro data: real GDP, hours worked, consumption, investment, real wages, prices, and the short-term nominal interest rate. CEE (2005) show that a version of the model estimated in this paper can replicate the impulse responses following a monetary policy shock identified in an unrestricted Vector Autoregression (VAR). As in Smets and Wouters (2003), the introduction of a larger number of shocks allows us to estimate the full model using the seven data series mentioned above. The marginal likelihood criterion, which captures the out-of-sample prediction performance, is used to test the NNS model against standard and Bayesian VAR models. We find that the NNS model has a fit comparable to that of Bayesian VAR models. These results are confirmed by a simple out-of-sample forecasting exercise. The restrictions implied by the NNS model lead to an improvement of the forecasting performance compared to standard VARs, in particular, at medium-term horizons. Bayesian NNS models, therefore, combine a sound, microfounded structure suitable for policy analysis with a good probabilistic description of the observed data and good forecasting performance.

Second, the introduction of a large number of frictions raises the question whether each of those frictions is really necessary to describe the seven data series. For example, CEE (2005) show that once one allows for nominal wage rigidity, there is no need for additional price rigidity in order to capture the impulse responses following a monetary policy shock. The Bayesian estimation methodology provides a natural framework for testing which frictions are empirically important by comparing the marginal likelihood of the various models. In contrast to CEE (2005), price and wage stickiness are found to be equally important. Indexation, on the other hand, is relatively unimportant in both goods and labor markets, confirming the single-equation results of Galí and Gertler (1999). While all the real frictions help in reducing the prediction errors of the NNS model, empirically the most important

are the investment adjustment costs. In the presence of wage stickiness, the introduction of variable capacity utilization is less important.

Finally, we use the estimated NNS model to address a number of key issues. First, what are the main driving forces of output developments in the United States? Broadly speaking, we confirm the analysis of Matthew D. Shapiro and Mark Watson (1988), who use a structural VAR methodology to examine the sources of business cycle fluctuations. While “demand” shocks such as the risk premium, exogenous spending, and investment-specific technology shocks explain a significant fraction of the short-run forecast variance in output, both wage mark-up (or labor supply) and, to a lesser extent, productivity shocks explain most of its variation in the medium to long run. Second, in line with Galí (1999) and Neville Francis and Valery A. Ramey (2004), productivity shocks have a significant short-run negative impact on hours worked. This is the case even in the flexible price economy, because of the slow adjustment of the two demand components following a positive productivity shock. Third, inflation developments are mostly driven by the price mark-up shocks in the short run and the wage mark-up shocks in the long run. Nevertheless, the model is able to capture the cross correlation between output and inflation at business cycle frequencies. Finally, in order to investigate the stability of the results, we estimate the NNS model for two subsamples: the “Great Inflation” period from 1966:2 to 1979:2 and the “Great Moderation” period from 1984:1 to 2004:4. We find that most of the structural parameters are stable over those two periods. The biggest difference concerns the variances of the structural shocks. In particular, the standard deviations of the productivity, monetary policy, and price mark-up shocks seem to have fallen in the second subsample, explaining the fall in the volatility of output growth and inflation in this period. We also detect a fall in the monetary policy response to output developments in the second subperiod.

In the next section, we discuss the linearized dynamic, stochastic, general-equilibrium (DSGE) model that is subsequently estimated. In Section II, the prior and posterior distribution of the structural parameters and the shock processes are discussed. In Section III, the model statistics and forecast performance are compared to those

of unconstrained VAR (and BVAR) models. In Section IV, the empirical importance of the different frictions are discussed. Finally, in Section V, we use the estimated model to discuss a number of key issues in business cycle analysis. Section VI contains the concluding remarks.

I. The Linearized DSGE Model

The DSGE model contains many frictions that affect both nominal and real decisions of households and firms. The model is based on CEE (2005) and Smets and Wouters (2003). As in Smets and Wouters (2005), we extend the model so that it is consistent with a balanced steady-state growth path driven by deterministic labor-augmenting technological progress. Households maximize a nonseparable utility function with two arguments (goods and labor effort) over an infinite life horizon. Consumption appears in the utility function relative to a time-varying external habit variable. Labor is differentiated by a union, so there is some monopoly power over wages, which results in an explicit wage equation and allows for the introduction of sticky nominal wages à la Guillermo A. Calvo (1983). Households rent capital services to firms and decide how much capital to accumulate given the capital adjustment costs they face. As the rental price of capital changes, the utilization of the capital stock can be adjusted at increasing cost. Firms produce differentiated goods, decide on labor and capital inputs, and set prices, again according to the Calvo model. The Calvo model in both wage and price setting is augmented by the assumption that prices that are not reoptimized are partially indexed to past inflation rates. Prices are therefore set in function of current and expected marginal costs, but are also determined by the past inflation rate. The marginal costs depend on wages and the rental rate of capital. Similarly, wages depend on past and expected future wages and inflation.

There are a few differences with respect to the model developed in Smets and Wouters (2005). First, the number of structural shocks is reduced to seven in order to match the number of observables that are used in estimation. Second, in both goods and labor markets we replace the Dixit-Stiglitz aggregator with an aggregator that allows for a time-varying demand elasticity, which depends on the relative price as in Kimball (1995). As shown by Eichenbaum and

Fischer (forthcoming), the introduction of this real rigidity allows us to estimate a more reasonable degree of price and wage stickiness.

In the rest of this section, we describe the log-linearized version of the DSGE model that we subsequently estimate using US data. All variables are log-linearized around their steady-state balanced growth path. Starred variables denote steady-state values.² We first describe the aggregate demand side of the model and then turn to the aggregate supply.

The aggregate resource constraint is given by

$$(1) \quad y_t = c_y c_t + i_y i_t + z_y z_t + \varepsilon_t^s.$$

Output (y_t) is absorbed by consumption (c_t), investment (i_t), capital-utilization costs that are a function of the capital utilization rate (z_t), and exogenous spending (ε_t^s); c_y is the steady-state share of consumption in output and equals $1 - g_y - i_y$, where g_y and i_y are respectively the steady-state exogenous spending-output ratio and investment-output ratio. The steady-state investment-output ratio in turn equals $(\gamma - 1 + \delta)k_y$, where γ is the steady-state growth rate, δ stands for the depreciation rate of capital, and k_y is the steady-state capital-output ratio. Finally, $z_y = R_*^k k_y$, where R_*^k is the steady-state rental rate of capital. We assume that exogenous spending follows a first-order autoregressive process with an IID-Normal error term and is also affected by the productivity shock as follows: $\varepsilon_t^s = \rho_g \varepsilon_{t-1}^s + \eta_t^s + \rho_{ga} \eta_t^a$. The latter is empirically motivated by the fact that, in estimation, exogenous spending also includes net exports, which may be affected by domestic productivity developments.

The dynamics of consumption follows from the consumption Euler equation and is given by

$$(2) \quad c_t = c_1 c_{t-1} + (1 - c_1) E_t c_{t+1} \\ + c_2 (l_t - E_t l_{t+1}) - c_3 (r_t - E_t \pi_{t+1} + \varepsilon_t^b),$$

where $c_1 = (\lambda\gamma)/(1 + \lambda\gamma)$, $c_2 = [(\sigma_c - 1)(W_*^h L_* / C_*)]/[\sigma_c(1 + \lambda\gamma)]$, and $c_3 = (1 - \lambda\gamma)/$

² Some details of the decisions faced by agents in the economy are given in the Model Appendix, available at http://www.e-aer.org/data/june07/20041254_app.pdf. An appendix with the full derivation of the steady state and the linearized model equations is available upon request.

$[(1 + \lambda\gamma)\sigma_c]$. Current consumption (c_t) depends on a weighted average of past and expected future consumption, and on expected growth in hours worked ($l_t - E_t l_{t+1}$), the ex ante real interest rate ($r_t - E_t \pi_{t+1}$), and a disturbance term ε_t^b . Under the assumption of no external habit formation ($\lambda = 0$) and log utility in consumption ($\sigma_c = 1$), $c_1 = c_2 = 0$ and the traditional purely forward-looking consumption equation is obtained. With steady-state growth, the growth rate γ marginally affects the reduced-form parameters in the linearized consumption equation. When the elasticity of intertemporal substitution (for constant labor) is smaller than one ($\sigma_c > 1$), consumption and hours worked are complements in utility and consumption depends positively on current hours worked and negatively on expected growth in hours worked (see Susanto Basu and Kimball 2002). Finally, the disturbance term ε_t^b represents a wedge between the interest rate controlled by the central bank and the return on assets held by the households. A positive shock to this wedge increases the required return on assets and reduces current consumption. At the same time, it also increases the cost of capital and reduces the value of capital and investment, as shown below.³ This shock has similar effects as so-called net-worth shocks in Ben S. Bernanke, Gertler, and Simon Gilchrist (1999) and Christiano, Roberto Motto, and Massimo Rostagno (2003), which explicitly model the external finance premium. The disturbance is assumed to follow a first-order autoregressive process with an IID-Normal error term: $\varepsilon_t^b = \rho_b \varepsilon_{t-1}^b + \eta_t^b$.

The dynamics of investment comes from the investment Euler equation and is given by

$$(3) \quad i_t = i_1 i_{t-1} + (1 - i_1) E_t i_{t+1} + i_2 q_t + \varepsilon_t^i,$$

where $i_1 = 1/(1 + \beta\gamma^{1-\sigma_c})$, $i_2 = [1/(1 + \beta\gamma^{1-\sigma_c})\gamma^2\varphi]$, φ is the steady-state elasticity of the capital adjustment cost function, and β is the discount factor applied by households. As in CEE (2005), a higher elasticity of the cost of adjusting capital reduces the sensitivity of investment (i_t) to

³ This latter effect makes this shock different from a discount factor shock (as in Smets and Wouters 2003), which affects only the intertemporal consumption Euler equation. In contrast to a discount factor shock, the risk premium shock helps to explain the comovement of consumption and investment.

the real value of the existing capital stock (q_t). Modeling capital adjustment costs as a function of the change in investment rather than its level introduces additional dynamics in the investment equation, which is useful in capturing the hump-shaped response of investment to various shocks. Finally, ε_t^i represents a disturbance to the investment-specific technology process and is assumed to follow a first-order autoregressive process with an IID-Normal error term: $\varepsilon_t^i = \rho_i \varepsilon_{t-1}^i + \eta_t^i$.

The corresponding arbitrage equation for the value of capital is given by

$$(4) \quad q_t = q_1 E_t q_{t+1} + (1 - q_1) E_t r_{t+1}^k - (r_t - *E_t \pi_{t+1} + \varepsilon_t^b),$$

where $q_1 = \beta\gamma^{-\sigma_c}(1 - \delta) = [(1 - \delta)/(R_*^k + (1 - \delta))]$. The current value of the capital stock (q_t) depends positively on its expected future value and the expected real rental rate on capital ($E_t r_{t+1}^k$) and negatively on the ex ante real interest rate and the risk premium disturbance.

Turning to the supply side, the aggregate production function is given by

$$(5) \quad y_t = \phi_p (\alpha k_t^s + (1 - \alpha) l_t + \varepsilon_t^a).$$

Output is produced using capital (k_t^s) and labor services (hours worked, l_t). Total factor productivity (ε_t^a) is assumed to follow a first-order autoregressive process: $\varepsilon_t^a = \rho_a \varepsilon_{t-1}^a + \eta_t^a$. The parameter α captures the share of capital in production, and the parameter ϕ_p is one plus the share of fixed costs in production, reflecting the presence of fixed costs in production.

As newly installed capital becomes effective only with a one-quarter lag, current capital services used in production (k_t^s) are a function of capital installed in the previous period (k_{t-1}) and the degree of capital utilization (z_t):

$$(6) \quad k_t^s = k_{t-1} + z_t.$$

Cost minimization by the households that provide capital services implies that the degree of capital utilization is a positive function of the rental rate of capital,

$$(7) \quad z_t = z_1 r_t^k,$$

where $z_1 = (1 - \psi)/\psi$ and ψ is a positive function of the elasticity of the capital utilization adjustment cost function and normalized to be between zero and one. When $\psi = 1$, it is extremely costly to change the utilization of capital and, as a result, the utilization of capital remains constant. In contrast, when $\psi = 0$, the marginal cost of changing the utilization of capital is constant and, as a result, in equilibrium the rental rate on capital is constant, as is clear from equation (7).

The accumulation of installed capital (k_t) is a function not only of the flow of investment but also of the relative efficiency of these investment expenditures as captured by the investment-specific technology disturbance

$$(8) \quad k_t = k_1 k_{t-1} + (1 - k_1) i_t + k_2 \varepsilon_t^i,$$

with $k_1 = (1 - \delta)/\gamma$ and $k_2 = (1 - (1 - \delta)/\gamma)(1 + \beta\gamma^{1-\sigma_c})\gamma^2\varphi$.

Turning to the monopolistic competitive goods market, cost minimization by firms implies that the price mark-up (μ_t^p), defined as the difference between the average price and the nominal marginal cost or the negative of the real marginal cost, is equal to the difference between the marginal product of labor (mpl_t) and the real wage (w_t):

$$(9) \quad \mu_t^p = mpl_t - w_t = \alpha(k_t^\alpha - l_t) + \varepsilon_t^a - w_t.$$

As implied by the second equality in (9), the marginal product of labor is itself a positive function of the capital-labor ratio and total factor productivity.

Due to price stickiness, as in Calvo (1983), and partial indexation to lagged inflation of those prices that can not be reoptimized, as in Smets and Wouters (2003), prices adjust only sluggishly to their desired mark-up. Profit maximization by price-setting firms gives rise to the following New-Keynesian Phillips curve:

$$(10) \quad \pi_t = \pi_1 \pi_{t-1} + \pi_2 E_t \pi_{t+1} - \pi_3 \mu_t^p + \varepsilon_t^p,$$

where $\pi_1 = \iota_p/(1 + \beta\gamma^{1-\sigma_c}\iota_p)$, $\pi_2 = \beta\gamma^{1-\sigma_c}/(1 + \beta\gamma^{1-\sigma_c}\iota_p)$, and $\pi_3 = 1/(1 + \beta\gamma^{1-\sigma_c}\iota_p)[(1 - \beta\gamma^{1-\sigma_c}\xi_p)(1 - \xi_p)/\xi_p((\phi_p - 1)\varepsilon_p + 1)]$. Inflation

(π_t) depends positively on past and expected future inflation, negatively on the current price mark-up, and positively on a price mark-up disturbance (ε_t^p). The price mark-up disturbance is assumed to follow an ARMA(1, 1) process: $\varepsilon_t^p = \rho_p \varepsilon_{t-1}^p + \eta_t^p - \mu_p \eta_{t-1}^p$, where η_t^p is an IID-Normal price mark-up shock. The inclusion of the MA term is designed to capture the high-frequency fluctuations in inflation.

When the degree of indexation to past inflation is zero ($\iota_p = 0$), equation (10) reverts to a standard, purely forward-looking Phillips curve ($\pi_1 = 0$). The assumption that all prices are indexed to either lagged inflation or the steady-state inflation rate ensures that the Phillips curve is vertical in the long run. The speed of adjustment to the desired mark-up depends, among others, on the degree of price stickiness (ξ_p), the curvature of the Kimball goods market aggregator (ε_p), and the steady-state mark-up, which in equilibrium is itself related to the share of fixed costs in production ($\phi_p - 1$) through a zero-profit condition. A higher ε_p slows down the speed of adjustment because it increases the strategic complementarity with other price setters. When all prices are flexible ($\xi_p = 0$) and the price-mark-up shock is zero, equation (10) reduces to the familiar condition that the price mark-up is constant, or equivalently that there are no fluctuations in the wedge between the marginal product of labor and the real wage.

Cost minimization by firms will also imply that the rental rate of capital is negatively related to the capital-labor ratio and positively to the real wage (both with unitary elasticity):

$$(11) \quad r_t^k = -(k_t - l_t) + w_t.$$

In analogy with the goods market, in the monopolistically competitive labor market, the wage mark-up will be equal to the difference between the real wage and the marginal rate of substitution between working and consuming (mrs_t),

$$(12) \quad \mu_t^w = w_t - mrs_t = w_t - \left(\sigma_l l_t + \frac{1}{1 - \lambda/\gamma} (c_t - \lambda \gamma c_{t-1}) \right),$$

where σ_l is the elasticity of labor supply with respect to the real wage and λ is the habit parameter in consumption.

Similarly, due to nominal wage stickiness and partial indexation of wages to inflation, real wages adjust only gradually to the desired wage mark-up:

$$(13) \quad w_t = w_1 w_{t-1} + (1 - w_1)(E_t w_{t+1} + E_t \pi_{t+1}) - w_2 \pi_t + w_3 \pi_{t-1} - w_4 \mu_t^w + \varepsilon_t^w,$$

with $w_1 = 1/(1 + \beta\gamma^{1-\sigma_c})$, $w_2 = (1 + \beta\gamma^{1-\sigma_c} \iota_w)/(1 + \beta\gamma^{1-\sigma_c})$, $w_3 = \iota_w/(1 + \beta\gamma^{1-\sigma_c})$, and $w_4 = 1/(1 + \beta\gamma^{1-\sigma_c})[(1 - \beta\gamma^{1-\sigma_c} \xi_w)(1 - \xi_w)/(\xi_w((\phi_w - 1)\varepsilon_w + 1))]$.

The real wage w_t is a function of expected and past real wages, expected, current, and past inflation, the wage mark-up, and a wage-markup disturbance (ε_t^w). If wages are perfectly flexible ($\xi_w = 0$), the real wage is a constant mark-up over the marginal rate of substitution between consumption and leisure. In general, the speed of adjustment to the desired wage mark-up depends on the degree of wage stickiness (ξ_w) and the demand elasticity for labor, which itself is a function of the steady-state labor market mark-up ($\phi_w - 1$) and the curvature of the Kimball labor market aggregator (ε_w). When wage indexation is zero ($\iota_w = 0$), real wages do not depend on lagged inflation ($w_3 = 0$). The wage-markup disturbance (ε_t^w) is assumed to follow an ARMA(1, 1) process with an IID-Normal error term: $\varepsilon_t^w = \rho_w \varepsilon_{t-1}^w + \eta_t^w - \mu_w \eta_{t-1}^w$. As in the case of the price mark-up shock, the inclusion of an MA term allows us to pick up some of the high-frequency fluctuations in wages.⁴

Finally, the model is closed by adding the following empirical monetary policy reaction function:

$$(14) \quad r_t = \rho r_{t-1} + (1 - \rho)\{r_\pi \pi_t + r_y(y_t - y_t^p)\} + r_{\Delta y}[(y_t - y_t^p) - (y_{t-1} - y_{t-1}^p)] + \varepsilon_t^r.$$

The monetary authorities follow a generalized Taylor rule by gradually adjusting the policy-controlled interest rate (r_t) in response to inflation and the output gap, defined as the difference between actual and potential output (John B. Taylor 1993). Consistently with the DSGE model, potential output is defined as the level of output that would prevail under flexible prices and wages in the absence of the two “mark-up” shocks.⁵ The parameter ρ captures the degree of interest rate smoothing. In addition, there is a short-run feedback from the change in the output gap. Finally, we assume that the monetary policy shocks (ε_t^r) follow a first-order autoregressive process with an IID-Normal error term: $\varepsilon_t^r = \rho_r \varepsilon_{t-1}^r + \eta_t^r$.

Equations (1) to (14) determine 14 endogenous variables: y_t , c_t , l_t , q_t , k_t^s , k_t , z_t , r_t^k , μ_t^p , π_t , μ_t^w , w_t , l_t , and r_t . The stochastic behavior of the system of linear rational expectations equations is driven by seven exogenous disturbances: total factor productivity (ε_t^a), investment-specific technology (ε_t^i), risk premium (ε_t^b), exogenous spending (ε_t^s), price mark-up (ε_t^p), wage mark-up (ε_t^w), and monetary policy (ε_t^r) shocks. Next we turn to the estimation of the model.

II. Parameter Estimates

The model presented in the previous section is estimated with Bayesian estimation techniques using seven key macroeconomic quarterly US time series as observable variables: the log difference of real GDP, real consumption, real investment and the real wage, log hours worked, the log difference of the GDP deflator, and the federal funds rate. A full description of the data used is given in the Data Appendix. The corresponding measurement equation is:

⁴ Alternatively, we could interpret this disturbance as a labor supply disturbance coming from changes in preferences for leisure.

⁵ In practical terms, we expand the model consisting of equations (1) to (14) with a flexible-price-and-wage version in order to calculate the model-consistent output gap. Note that the assumption of treating the wage equation disturbance as a wage mark-up disturbance rather than a labor supply disturbance coming from changed preferences has implications for our calculation of potential output.

$$(15) \quad Y_t = \begin{bmatrix} dlGDP_t \\ dlCONS_t \\ dlINV_t \\ dlWAG_t \\ lHOURS_t \\ dlP_t \\ FEDFUNDS_t \end{bmatrix} = \begin{bmatrix} \bar{\gamma} \\ \bar{\gamma} \\ \bar{\gamma} \\ \bar{\gamma} \\ \bar{l} \\ \bar{\pi} \\ \bar{r} \end{bmatrix} + \begin{bmatrix} y_t - y_{t-1} \\ c_t - c_{t-1} \\ i_t - i_{t-1} \\ w_t - w_{t-1} \\ l_t \\ \pi_t \\ r_t \end{bmatrix},$$

where l and dl stand for 100 times log and log difference, respectively; $\bar{\gamma} = 100(\gamma - 1)$ is the common quarterly trend growth rate to real GDP, consumption, investment and wages; $\bar{\pi} = 100(\Pi_* - 1)$ is the quarterly steady-state inflation rate; and $\bar{r} = 100(\beta^{-1}\gamma^\sigma\Pi_* - 1)$ is the steady-state nominal interest rate. Given the estimates of the trend growth rate and the steady-state inflation rate, the latter will be determined by the estimated discount rate. Finally, \bar{l} is steady-state hours worked, which is normalized to be equal to zero.

First, we estimate the mode of the posterior distribution by maximizing the log posterior function, which combines the prior information on the parameters with the likelihood of the data. In a second step, the Metropolis-Hastings algorithm is used to get a complete picture of the posterior distribution and to evaluate the marginal likelihood of the model.⁶ The model is

⁶ See Smets and Wouters (2003) for a more elaborate description of the methodology. All estimations are done with Dynare (<http://www.cpremap.cnrs.fr/dynare>). A sample of 250,000 draws was created (neglecting the first 50,000 draws). The Hessian resulting from the optimization procedure was used for defining the transition probability function that generates the new proposed draw. A step size of 0.3 resulted in a rejection rate of 0.65. The resulting sample properties are not sensitive to the step size. Two methods were used to test the stability of the sample. The first convergence diagnostic is based on Stephen P. Brooks and Andrew Gelman (1998) and compares between and within moments of multiple chains. These tests are implemented in Dynare. The second method to evaluate the stability is a graphical test based on the cumulative mean minus the overall mean (see Luc Bauwens, Michel Lubrano, and Jean-François Richards 2000). An exact statistical test

estimated over the full sample period from 1966:1 to 2004:4. In Section VD, we estimate the model over two subperiods (1966:1–1979:2 and 1984:1–2004:4) in order to investigate the stability of the estimated parameters.⁷

A. Prior Distribution of the Parameters

The priors on the stochastic processes are harmonized as much as possible. The standard errors of the innovations are assumed to follow an inverse-gamma distribution with a mean of 0.10 and two degrees of freedom, which corresponds to a rather loose prior. The persistence of the AR(1) processes is beta distributed with mean 0.5 and standard deviation 0.2. A similar distribution is assumed for the MA parameter in the process for the price and wage mark-up. The quarterly trend growth rate is assumed to be Normal distributed with mean 0.4 (quarterly growth rate) and standard deviation 0.1. The steady-state inflation rate and the discount rate are assumed to follow a gamma distribution with a mean of 2.5 percent and 1 percent on an annual basis.

Five parameters are fixed in the estimation procedure. The depreciation rate δ is fixed at 0.025 (on a quarterly basis) and the exogenous spending-GDP ratio g_y is set at 18 percent. Both these parameters would be difficult to estimate unless the investment and exogenous spending ratios were used directly in the measurement equation. Three other parameters are clearly not identified: the steady-state mark-up in the labor market (λ_w), which is set at 1.5, and the curvature parameters of the Kimball aggregators in the goods and labor market (ϵ_p and ϵ_w), which are both set at 10.

for the stability of the sample is complicated by the highly autocorrelated nature of the MH-sampler. From an economic point of view, however, the differences between subsamples and independent samples of size 100,000 or more are negligible.

⁷ The dataset used generally starts in 1947. In previous versions of this paper, however, we found that the first ten years are not representative of the rest of the sample, so that we decided to shorten the sample to 1957:1–2004:4. In addition, below in Section IV we use the first ten years as a training sample for calculating the marginal likelihood of unconstrained VARs, so that the effective sample starts in 1966:1.

TABLE 1A—PRIOR AND POSTERIOR DISTRIBUTION OF STRUCTURAL PARAMETERS

	Prior distribution			Posterior distribution			
	Distr.	Mean	St. Dev.	Mode	Mean	5 percent	95 percent
φ	Normal	4.00	1.50	5.48	5.74	3.97	7.42
σ_c	Normal	1.50	0.37	1.39	1.38	1.16	1.59
h	Beta	0.70	0.10	0.71	0.71	0.64	0.78
ξ_w	Beta	0.50	0.10	0.73	0.70	0.60	0.81
σ_l	Normal	2.00	0.75	1.92	1.83	0.91	2.78
ξ_p	Beta	0.50	0.10	0.65	0.66	0.56	0.74
ι_w	Beta	0.50	0.15	0.59	0.58	0.38	0.78
ι_p	Beta	0.50	0.15	0.22	0.24	0.10	0.38
ψ	Beta	0.50	0.15	0.54	0.54	0.36	0.72
Φ	Normal	1.25	0.12	1.61	1.60	1.48	1.73
r_π	Normal	1.50	0.25	2.03	2.04	1.74	2.33
ρ	Beta	0.75	0.10	0.81	0.81	0.77	0.85
r_y	Normal	0.12	0.05	0.08	0.08	0.05	0.12
$r_{\Delta y}$	Normal	0.12	0.05	0.22	0.22	0.18	0.27
$\bar{\pi}$	Gamma	0.62	0.10	0.81	0.78	0.61	0.96
$100(\beta^{-1} - 1)$	Gamma	0.25	0.10	0.16	0.16	0.07	0.26
\bar{l}	Normal	0.00	2.00	-0.1	0.53	-1.3	2.32
$\bar{\gamma}$	Normal	0.40	0.10	0.43	0.43	0.40	0.45
α	Normal	0.30	0.05	0.19	0.19	0.16	0.21

Note: The posterior distribution is obtained using the Metropolis-Hastings algorithm.

The parameters describing the monetary policy rule are based on a standard Taylor rule: the long-run reaction on inflation and the output gap are described by a Normal distribution with mean 1.5 and 0.125 (0.5 divided by 4) and standard errors 0.125 and 0.05, respectively. The persistence of the policy rule is determined by the coefficient on the lagged interest rate rate, which is assumed to be Normal around a mean of 0.75 with a standard error of 0.1. The prior on the short-run reaction coefficient to the change in the output gap is 0.125.

The parameters of the utility function are assumed to be distributed as follows. The intertemporal elasticity of substitution is set at 1.5 with a standard error of 0.375; the habit parameter is assumed to fluctuate around 0.7 with a standard error of 0.1, and the elasticity of labor supply is assumed to be around 2 with a standard error of 0.75. These are all quite standard calibrations. The prior on the adjustment cost parameter for investment is set around 4 with a standard error of 1.5 (based on CEE 2005) and the capacity utilization elasticity is set at 0.5 with a standard error of 0.15. The share of fixed costs in the production function is assumed to have a prior mean of 0.25. Finally, there are the parameters describing the price and wage setting. The Calvo probabilities are assumed to be

around 0.5 for both prices and wages, suggesting an average length of price and wage contracts of half a year. This is compatible with the findings of Mark Bils and Peter J. Klenow (2004) for prices. The prior mean of the degree of indexation to past inflation is also set at 0.5 in both goods and labor markets.⁸

B. Posterior Estimates of the Parameters

Table 1 gives the mode, the mean, and the 5 and 95 percentiles of the posterior distribution of the parameters obtained by the Metropolis-Hastings algorithm.

The trend growth rate is estimated to be around 0.43, which is somewhat smaller than the average growth rate of output per capita over the sample. The posterior mean of the steady-state inflation rate over the full sample is about 3 percent on an annual basis. The mean of the discount rate is estimated to be quite small (0.65 percent on an annual basis). The implied mean steady-state nominal and real interest

⁸ We have analyzed the sensitivity of the estimation results to the prior assumptions by increasing the standard errors of the prior distributions of the behavioral parameters by 50 percent. Overall, the estimation results are very similar.

TABLE 1B—PRIOR AND POSTERIOR DISTRIBUTION OF SHOCK PROCESSES

	Prior distribution			Posterior distribution			
	Distr.	Mean	St. Dev.	Mode	Mean	95 percent	5 percent
σ_a	Invgamma	0.10	2.00	0.45	0.45	0.41	0.50
σ_b	Invgamma	0.10	2.00	0.24	0.23	0.19	0.27
σ_g	Invgamma	0.10	2.00	0.52	0.53	0.48	0.58
σ_l	Invgamma	0.10	2.00	0.45	0.45	0.37	0.53
σ_r	Invgamma	0.10	2.00	0.24	0.24	0.22	0.27
σ_p	Invgamma	0.10	2.00	0.14	0.14	0.11	0.16
σ_w	Invgamma	0.10	2.00	0.24	0.24	0.20	0.28
ρ_a	Beta	0.50	0.20	0.95	0.95	0.94	0.97
ρ_b	Beta	0.50	0.20	0.18	0.22	0.07	0.36
ρ_g	Beta	0.50	0.20	0.97	0.97	0.96	0.99
ρ_l	Beta	0.50	0.20	0.71	0.71	0.61	0.80
ρ_r	Beta	0.50	0.20	0.12	0.15	0.04	0.24
ρ_p	Beta	0.50	0.20	0.90	0.89	0.80	0.96
ρ_w	Beta	0.50	0.20	0.97	0.96	0.94	0.99
μ_p	Beta	0.50	0.20	0.74	0.69	0.54	0.85
μ_w	Beta	0.50	0.20	0.88	0.84	0.75	0.93
ρ_{ga}	Beta	0.50	0.20	0.52	0.52	0.37	0.66

Note: The posterior distribution is obtained using the Metropolis-Hastings algorithm.

rates are, respectively, about 6 percent and 3 percent on an annual basis.

A number of observations are worth making regarding the estimated processes for the exogenous shock variables (Table 1B). Overall, the data appear to be very informative on the stochastic processes for the exogenous disturbances. The productivity, the government spending, and the wage mark-up processes are estimated to be the most persistent, with an AR(1) coefficient of 0.95, 0.97, and 0.96, respectively. The mean of the standard error of the shock to the productivity process is 0.45. The high persistence of the productivity and wage mark-up processes implies that, at long horizons, most of the forecast error variance of the real variables will be explained by those two shocks. In contrast, both the persistence and the standard deviation of the risk premium and monetary policy shock are relatively low (0.18 and 0.12, respectively).

Turning to the estimates of the main behavioral parameters, it turns out that the mean of the posterior distribution is typically relatively close to the mean of the prior assumptions. There are a few notable exceptions. The degree of both price and wage stickiness is estimated to be quite a bit higher than 0.5. The average duration of wage contracts is somewhat less than a year; whereas the average duration of price contracts is about three quarters. The

mean of the degree of price indexation (0.24) is, on the other hand, estimated to be much less than 0.5.⁹ In addition, the elasticity of the cost of changing investment is estimated to be higher than assumed a priori, suggesting an even slower response of investment to changes in the value of capital. Finally, the posterior mean of the fixed cost parameter is estimated to be much higher than assumed in the prior distribution (1.6) and the share of capital in production is estimated to be much lower (0.19). Overall, it appears that the data are quite informative on the behavioral parameters, as indicated by the lower variance of the posterior distribution relative to the prior distribution. Two exceptions are the elasticity of labor supply and the elasticity of the cost of changing the utilization of capital, where the posterior and prior distributions are quite similar.¹⁰

Finally, turning to the monetary policy reaction function parameters, the mean of the long-run reaction coefficient to inflation is estimated to be relatively high (2.0). There is a considerable degree of interest rate smoothing, as the

⁹ When relaxing the prior distributions, it turns out that the degree of wage stickiness rises even more, whereas the degree of price indexation falls by more.

¹⁰ Figures with the prior and posterior distributions of all the parameters are available upon request.

mean of the coefficient on the lagged interest rate is estimated to be 0.81. Policy does not appear to react very strongly to the output gap level (0.09), but does respond strongly to changes in the output-gap (0.22) in the short run.

III. Forecast Performance: Comparison with VAR Models

In this section, we compare the out-of-sample forecast performance of the estimated DSGE model with that of various VARs estimated on the same dataset. The marginal likelihood, which can be interpreted as a summary statistic for the model's out-of-sample prediction performance, forms a natural benchmark for comparing the DSGE model with alternative specifications and other statistical models.¹¹ As Sims (2003) has pointed out, however, it is important to use a training sample in order to standardize the prior distribution across widely different models. In order to check for robustness, we also consider a more traditional out-of-sample root mean squared error (RMSE) forecast exercise in this section.

Table 2 compares the marginal likelihood of the DSGE model and various unconstrained VAR models, all estimated over the full sample period (1966:1–2004:4) and using the period 1956:1–1965:4 as a training sample. Several results are worth emphasizing. First, the tightly parameterized DSGE model performs much better than an unconstrained VAR in the same vector of observable variables, Y_t (first column of Table 2). The bad empirical performance of unconstrained VARs may not be too surprising, as it is known that overparameterized models typically perform poorly in out-of-sample fore-

TABLE 2—COMPARISON OF THE MARGINAL LIKELIHOOD OF ALTERNATIVE VAR MODELS AND THE DSGE MODEL

Order of the VAR	No other prior	Sims and Zha (1998) prior
VAR(1)	−928.0	−940.9
VAR(2)	−966.6	−915.8
VAR(3)	−1018.1	−908.7
VAR(4)	−1131.2	−906.6
VAR(5)	—	−907.7
Memo: DSGE model	−905.8	−905.8

Note: In order to increase the comparability of the marginal likelihood of the various models, all models are estimated using the period 1956:1–1965:4 as a training sample (Sims 2003).

cast exercises. One indication of this is that the marginal likelihood of the unconstrained VAR model deteriorates quickly as the lag order increases. For that reason, in the second column of Table 2, we consider the Bayesian VAR model proposed by Christopher A. Sims and Tao Zha (1998). This BVAR combines a Minnesota-type prior (see Robert B. Litterman 1984) with priors that take into account the degree of persistence and cointegration in the variables. In order to allow the data to decide on the degree of persistence and cointegration, in this BVAR we enter real GDP, consumption, investment, and the real wage in log levels. When setting the tightness of the prior, we choose a set of parameters recommended by Sims (2003) for quarterly data.¹² The second column of Table 2 shows that the marginal likelihood of the Sims-Zha BVAR increases significantly compared to the unconstrained VAR. Moreover, the best BVAR model (BVAR(4)) does as well as the DSGE model.¹³

Overall, the comparison of marginal likelihoods

¹¹ As discussed in John Geweke (1998), the Metropolis-Hastings-based sample of the posterior distribution can be used to evaluate the marginal likelihood of the model. Following Geweke (1998), we calculate the modified harmonic mean to evaluate the integral over the posterior sample. An alternative approximation is the Laplace approximation around the posterior mode, which is based on a normal distribution. In our experience, the results of both approximations are very close in the case of our estimated DSGE model. This is not too surprising, given the generally close correspondence between the histograms of the posterior sample and the normal distribution around the estimated mode for the individual parameters. Given the large advantage of the Laplace approximation in terms of computational costs, we will use this approximation for comparing alternative model specifications in the next section.

¹² In order to determine the tightness of the priors, we use standard values as suggested in Sims (2003) (see also Sims and Zha 1998). In particular, the decay parameter is set at 1.0, the overall tightness is set at 10, the parameter determining the weight on the “sum of coefficients” or “own-persistence” is set at 2.0, and the parameter determining the weight on the “co-persistence” is set at 5. Moreover, the vector of prior standard deviations of the equation shocks is based on the VAR(1) residuals estimated over the training period.

¹³ The marginal likelihood can be further increased by optimizing the tightness of the “own persistence” prior. For example, setting this parameter equal to 10 increases the marginal likelihood of the BVAR(4) model to −896.

TABLE 3—OUT-OF-SAMPLE PREDICTION PERFORMANCE

	GDP	dP	Fedfunds	Hours	Wage	CONS	INV	Overall
<i>VAR(1)</i>	<i>RMSE-statistic for different forecast horizons</i>							
1q	0.60	0.25	0.10	0.46	0.64	0.60	1.62	-12.87
2q	0.94	0.27	0.18	0.78	1.02	0.95	2.96	-8.19
4q	1.64	0.34	0.36	1.45	1.67	1.54	5.67	-3.25
8q	2.40	0.53	0.64	2.13	2.88	2.27	8.91	1.47
12q	2.78	0.63	0.79	2.41	4.09	2.74	10.97	2.36
<i>BVAR(4)</i>	<i>Percentage gains (+) or losses (-) relative to VAR(1) model</i>							
1q	2.05	14.14	-1.37	-3.43	2.69	12.12	2.54	3.25
2q	-2.12	15.15	-16.38	-7.32	-0.29	10.07	2.42	0.17
4q	-7.21	31.42	-12.61	-8.58	-3.82	1.42	0.43	0.51
8q	-15.82	33.36	-13.26	-13.94	-8.98	-8.19	-11.58	-4.10
12q	-15.55	37.59	-13.56	-4.66	-15.87	-3.10	-23.49	-9.84
<i>DSG</i>	<i>Percentage gains (+) or losses (-) relative to VAR(1) model</i>							
1q	5.68	2.05	-8.24	0.68	5.99	20.16	9.22	3.06
2q	14.93	10.62	-17.22	10.34	6.20	25.85	16.79	2.82
4q	20.17	46.21	1.59	19.52	9.21	26.18	21.42	6.82
8q	22.55	68.15	28.33	22.34	15.72	21.82	25.95	11.50
12q	32.17	74.15	40.32	27.05	21.88	23.28	41.61	13.51

Notes: All models are estimated starting in 1966:1. The forecast period is 1990:1–2004:4. VAR(1) and BVAR(4) models are reestimated each quarter, the DSGE model each year. The overall measure of forecast performance is the log determinant of the uncentered forecast error covariance matrix. Gains and losses in the overall measure are expressed as the difference in the overall measure divided by the number of variables and by two to convert the variance to standard errors (times 100).

shows that the estimated DSGE model can compete with standard BVAR models in terms of empirical one-step-ahead prediction performance. These results are confirmed by a more traditional out-of-sample forecasting exercise reported in Table 3. Table 3 reports out-of-sample RMSEs for different forecast horizons over the period 1990:1 to 2004:4. For this exercise, the VAR(1), BVAR(4), and DSGE model were initially estimated over the sample 1966:1–1989:4. The models were then used to forecast the seven data series contained in Y_t from 1990:1 to 2004:4, whereby the VAR(1) and BVAR(4) models were reestimated every quarter, and the DSGE model was reestimated every year. The measure of overall performance reported in the last column of Table 3 is the log determinant of the uncentered forecast error covariance matrix.

The out-of-sample forecast statistics confirm the good forecast performance of the DSGE model relative to the VAR and BVAR models. At the one-quarter-ahead horizon, the BVAR(4) and the DSGE model improve with about the same magnitude over the VAR(1) model, confirming the results from Table 2. Over longer horizons up to three years, however, the DSGE model does considerably better than both the VAR(1) and BVAR(4) model. Somewhat sur-

prisingly, the BVAR(4) model performs worse than the simple VAR(1) model at longer horizons. Moreover, the improvement appears to be quite uniform across the seven macro variables.

IV. Model Sensitivity: Which Frictions Are Empirically Important?

The introduction of a large number of frictions raises the question of which of those are really necessary to capture the dynamics of the data. In this section, we examine the contribution of each of the frictions to the marginal likelihood of the DSGE model.

Table 4 presents the estimates of the mode of the parameters and the marginal likelihood when each friction (price and wage stickiness, price and wage indexation, investment adjustment costs and habit formation, capital utilization, and fixed costs in production) is drastically reduced one at a time. This table also gives an idea of the robustness of the parameters and the model performance with respect to the various frictions included in the model. For comparison, the first column reproduces the baseline estimates (mode of the posterior) and the marginal likelihood based on the Laplace approximation for the model without training sample.

TABLE 4—TESTING THE EMPIRICAL IMPORTANCE OF THE NOMINAL AND REAL FRICTIONS IN THE DSGE MODEL

Base	$\xi_p = 0.1$	$\xi_w = 0.1$	$\iota_p = 0.0$	$\iota_w = 0.0$	$\varphi = 0.1$	$h = 0.1$	$\psi = 0.99$	$\Phi = 1.1$	
<i>Marginal likelihood</i>									
	-923	-975	-973	-918	-927	-1084	-959	-924	-949
<i>Mode of the structural parameters</i>									
φ	5.48	4.41	2.78	5.45	5.62	0.10	1.26	5.33	5.19
σ_c	1.39	1.31	1.80	1.43	1.42	2.78	1.90	1.39	1.27
h	0.71	0.70	0.34	0.70	0.71	0.12	0.10	0.70	0.71
ξ_w	0.73	0.55	0.10	0.75	0.75	0.89	0.73	0.73	0.78
σ_l	1.92	1.48	0.25	1.91	1.91	5.24	1.21	1.79	2.33
ξ_p	0.65	0.10	0.48	0.66	0.69	0.86	0.62	0.59	0.80
ι_w	0.59	0.71	0.68	0.61	0.01	0.39	0.61	0.63	0.58
ι_p	0.22	0.84	0.24	0.01	0.24	0.08	0.21	0.21	0.19
ψ	0.54	0.82	0.66	0.54	0.50	0.02	0.69	0.99	0.45
Φ	1.61	1.79	1.64	1.60	1.61	1.15	1.44	1.62	1.10
r_π	2.03	2.15	2.15	2.01	2.01	2.03	2.24	2.04	1.98
r	0.81	0.79	0.75	0.81	0.82	0.84	0.81	0.80	0.80
r_y	0.08	0.08	0.08	0.08	0.09	0.23	0.12	0.08	0.10
$r_{\Delta y}$	0.22	0.21	0.25	0.22	0.22	0.30	0.29	0.23	0.25
α	0.19	0.21	0.20	0.19	0.19	0.20	0.19	0.18	0.13
<i>Mode of the autoregressive parameters of the exogenous shock processes</i>									
ρ_a	0.95	0.96	0.97	0.96	0.95	0.99	0.97	0.96	0.96
ρ_b	0.18	0.19	0.67	0.18	0.18	0.89	0.79	0.18	0.28
ρ_g	0.97	0.96	0.97	0.97	0.97	0.99	0.97	0.97	0.96
ρ_l	0.71	0.71	0.78	0.70	0.69	0.99	0.90	0.73	0.74
ρ_r	0.12	0.14	0.13	0.12	0.11	0.02	0.03	0.13	0.11
ρ_p	0.90	0.97	0.94	0.88	0.88	0.60	0.93	0.92	0.85
ρ_w	0.97	0.98	0.98	0.97	0.97	0.92	0.98	0.97	0.95
μ_p	0.74	0.20	0.71	0.59	0.77	0.34	0.76	0.71	0.67
μ_w	0.88	0.75	0.14	0.91	0.88	0.96	0.95	0.90	0.87

We focus first on the nominal frictions. Reducing the degree of nominal price and wage stickiness to a Calvo probability of 0.10 is about equally costly in terms of a deterioration of the marginal likelihood. In both cases the marginal likelihood falls very significantly by about 50. A lower degree of price stickiness leads to a strong increase in the estimated degree of price indexation from 0.22 to 0.84. In addition, the variance and the persistence of the price mark-up shocks increase as a result. The other parameters are less affected. The main impact of reducing the degree of wage stickiness on the other parameters concerns the elasticity of wages with respect to employment: the labor supply elasticity becomes much smaller and falls from a value of 1.92 to 0.25. In terms of short-run dynamics, these changes more or less cancel out, leaving the impact of labor effort on wage dynamics unaffected. In this case, the variance and the persistence of the wage mark-up shock increases.

While both Calvo frictions are empirically quite important, neither price nor wage index-

ation plays a very important role in the model dynamics. On the contrary, restricting the price indexation parameter to a very low value of 0.01 leads to an improvement of the marginal likelihood, suggesting that empirically it would be better to leave this friction out. Moreover, leaving out either friction does not have any noticeable impact on the other parameters.

Turning to the real frictions, the most important in terms of the marginal likelihood are the investment adjustment costs. Reducing the elasticity of adjustment costs to a very low level leads to a deterioration of the marginal likelihood by 160. Also, reducing habit formation in consumption is quite costly, although much less so than reducing investment adjustment costs. The reduced hump-shaped endogenous dynamics of the model due to these restrictions is compensated mainly by higher and more persistent exogenous shocks to productivity, investment, consumption, and government spending. The other real frictions fall, while the nominal rigidities increase. The presence of variable capital

utilization does not seem to matter for the model's performance. Shutting this off comes at no cost. What is costly is to reduce the share of fixed costs in production to 10 percent. Contrary to the discussion in Robert G. King and Sergio T. Rebelo (2000), the absence of variable capital utilization does not increase the standard error of the productivity shock in our model. In contrast, reducing the fixed costs in production does mechanically increase the standard deviation of the productivity shock.

Overall, the results from this sensitivity exercise illustrate that the estimated parameters appear relatively robust to changes in the frictions, one by one. Price and wage indexation and variable capital utilization are of minor importance in terms of the overall empirical performance of the model. On the real side, investment adjustment costs are the most important friction. On the nominal side, both wage and price stickiness are very important.

V. Applications

After having shown that the estimated model fits the US macroeconomic data quite well, we use it to investigate a number of key macroeconomic issues. In this section, we address the following questions. First, what are the main driving forces of output? Second, can the model replicate the cross correlation between output and inflation? Third, what is the effect of a productivity shock on hours worked? And fourth, why have output and inflation become less volatile? We study these issues in each subsection in turn.

A. What Are the Main Driving Forces of Output?

Figure 1 gives the forecast error variance decomposition of output, inflation, and the federal funds rate at various horizons based on the mode of the model's posterior distribution reported in Section III. In the short run (within a year) movements in real GDP are primarily driven by the exogenous spending shock and the two shocks that affect the intertemporal Euler equations, i.e., the risk premium shock which affects both the consumption and investment Euler equation and the investment-specific technology shock which affects the investment Euler equation. Together, they account for more




FIGURE 1. FORECAST ERROR VARIANCE DECOMPOSITION
(At the mode of the posterior distribution)

than 50 percent of the forecast error variance of output up to one year. Each of those shocks can be categorized as “demand” shocks in the sense that they have a positive effect on output, hours worked, inflation, and the nominal interest rate under the estimated policy rule. This is illustrated in Figure 2, which shows the estimated mean impulse response functions to each of those three shocks. Not surprisingly, the risk premium shock explains a big part of the short-run variations in consumption, while the investment shock explains the largest part of investment in the short run (not shown).¹⁴

In line with the results of Matthew D. Shapiro and Mark Watson (1989), however, it is primarily two “supply” shocks, the productivity and the wage mark-up shock, that account for most of the output variations in the medium to long run. Indeed, even at the two-year horizon, to-

¹⁴ The full set of impulse response functions, as well as the associated confidence sets, are available upon request.




FIGURE 2. THE ESTIMATED MEAN IMPULSE RESPONSES TO "DEMAND" SHOCKS

Note: Bold solid line: risk premium shock; thin solid line: exogenous spending shock; dashed line: investment shock.

gether the two shocks account for more than 50 percent of the variations in output. In the longer run, the wage mark-up shock dominates the productivity shock. Those shocks also become dominant forces in the long-run developments of consumption and, to a lesser extent, investment. Not surprisingly, the wage-markup shock is also the dominant factor behind long-run movements in hours worked. As shown in Figure 3, a typical positive wage mark-up shock gradually reduces output and hours worked by 0.8 and 0.6 percent, respectively. Confirming the large identified VAR literature on the role of monetary policy shocks (e.g. Christiano, Eichenbaum, and Evans 2000), monetary policy shocks contribute only a small fraction of the forecast variance of output at all horizons.

Figure 4 shows the historical contribution of each of four types of shocks (productivity, demand, monetary policy, and mark-up shocks) to annual output growth over the sample period. It is interesting to compare the main sources of the various recessions over this period. While the recessions of the early 1990s and the beginning of the new millennium are driven mainly by demand shocks, the recession of 1974 is due primarily to positive mark-up shocks (associated with the oil crisis). Monetary policy shocks play a dominant role only in the recession of the early 1980s when the Federal Reserve, under the chairmanship of Paul Volker, started the disinflation process.




FIGURE 3. THE ESTIMATED IMPULSE RESPONSE TO A WAGE MARK-UP SHOCK

Note: The solid line is the mean impulse response; the dotted lines are the 10 percent and 90 percent posterior intervals.

B. Determinants of Inflation and the Output-Inflation Cross Correlation

Figure 1 also contains the variance decomposition of inflation. It is quite clear that, at all horizons, price and wage mark-ups are the most important drivers of inflation. In the short run, price mark-ups dominate, whereas in the medium to long run, wage mark-ups become relatively more important. Even at the medium-to long-run horizons, the other shocks explain only a minor fraction of the total variation in inflation. Similarly, monetary policy shocks account for only a small fraction of inflation volatility. This is also clear from Figure 4, which depicts the historical contribution of the different types of shocks to inflation over the sample period. The dominant source of secular shifts in inflation is driven by price and wage mark-up shocks. Monetary policy did, however, play a role in the rise of inflation in the 1970s and the subsequent disinflation during the Volker period. Moreover, negative demand shocks contributed to low inflation in the early 1990s and the start of the new millennium.

There are at least two reasons why the various demand and productivity shocks have only limited effects on inflation. First, the estimated slope of the New Keynesian Phillips curve is very small, so that only large and persistent changes in the marginal cost will have an impact on inflation. Second, and more importantly,




FIGURE 4. HISTORICAL DECOMPOSITION OF GDP GROWTH AND INFLATION
 (Annual per capita GDP growth (deviation from trend growth))

Notes: The demand shocks include the risk premium, investment-specific technology, and exogenous spending shocks; the mark-up shocks include the price and wage mark-up shocks. Trend per-capita growth is estimated at 1.73 percent, whereas mean inflation is estimated at 3.17 percent.




FIGURE 5. THE ACTUAL AND MODEL-BASED CROSS-CORRELATION FUNCTION BETWEEN OUTPUT AND INFLATION
Note: Output is Hodrick-Prescott filtered real GDP.

under the estimated monetary policy reaction function, the Fed responds quite aggressively to emerging output gaps and their impact on inflation. This is reflected in the fact that at the short- and medium-term horizon more than 60 percent of variations in the nominal interest rate are due to the various demand and productivity shocks, in particular the risk premium shock (third panel of Figure 1). Only in the long run does the wage mark-up shock become a dominant source of movements in nominal interest rates.

In the light of these results, it is interesting to see to what extent our model can replicate the empirical correlation function between output and inflation as, for example, highlighted in Galí and Gertler (1999). Figure 5 plots the empirical correlation function of output (detrended using the Hodrick-Prescott filter) and inflation (estimated over the period 1966:1–2004:4), as well as the median and the 5 percent and 95 percent equivalent generated by the model's posterior distribution. In order to generate this distribution, 1,000 draws from the posterior distribution of the model parameters are used to generate artificial samples of output and inflation of the same sample size as the actual dataset. For each of those 1,000 artificial samples, the autocorrelation function is calculated and the median and 5 and 95 percentiles are derived. Figure 5 clearly shows that the DSGE model is able to replicate both the negative correlation between inflation one to two years in the past




FIGURE 6. THE IMPULSE RESPONSES TO A MONETARY POLICY SHOCK

Note: The solid line is the mean impulse response; the dotted lines are the 10 percent and 90 percent posterior intervals.

and current output, and the positive correlation between current output and inflation one year ahead. Moreover, the correlations generated by the DSGE model are significantly different from zero. Decomposing the cross-covariance function in contributions by the different types of shocks, we find that the negative correlation between current inflation and future output is driven primarily by the price and wage mark-up shocks. In contrast, the positive correlation between the current output gap and future inflation is the result of both demand shocks and mark-up shocks. Monetary policy shocks do not play a role for two reasons. First, they account for only a small fraction of inflation and output developments. Second, as shown in Figure 6, according to the estimated DSGE model, the peak effect of a policy shock on inflation occurs before its peak effect on output.

C. The Effect of a Productivity Shock on Hours Worked

Following Galí (1999), there has been a lively debate about the effects of productivity shocks on hours worked and about the implications of this finding for the role of those shocks in US business cycles. Galí (1999), Francis and Ramey (2005), and Galí and Pau Rabanal (2004) have argued that due to the presence of nominal price rigidities, habit formation, and




FIGURE 7. THE ESTIMATED IMPULSE RESPONSES TO A PRODUCTIVITY SHOCK

Note: The solid lines represent the estimated actual mean responses and the 10 percent and 90 percent posterior interval; the dashed lines represent the counterfactual flexible-wage-and-price responses.

adjustment costs to investment, positive productivity shocks lead to an immediate fall in hours worked. Given the strongly positive correlation between output and hours worked over the business cycle, this implies that productivity shocks cannot play an important role in the business cycle. In contrast, using alternative VAR specifications and identification strategies, Christiano, Eichenbaum, and Robert Vigfusson (2004), Luca Dedola and Stefano Neri (2004), and Gert Peersman and Roland Straub (2005) have argued that the empirical evidence on the effect of a productivity shock on hours worked is not very robust and could be consistent with a positive impact on hours worked.

In Section VA we have already discussed that productivity shocks play an important, but not dominant, role in driving output developments beyond the one-year horizon in our estimated model. At business cycle frequencies, they account for about 25–30 percent of the forecast error variance. Figure 7 presents the response of the actual and the flexible-price level of output, hours worked, and nominal interest rate to a productivity shock in the estimated model. Overall, the estimates confirm the analysis of Galí (1999) and Francis and Ramey (2004). A positive productivity shock leads to an expansion of aggregate demand, output, and real wages, but an immediate and significant reduc-

tion in hours worked. Hours worked turn significantly positive only after two years.¹⁵ Under the estimated monetary policy reaction function, nominal and real interest rates fall, but not enough to prevent the opening up of an output gap and a fall in inflation. Moreover, our estimation results show that it is mainly the estimated degree of habit persistence and the importance of capital adjustment costs that explain the negative impact of productivity on hours worked, thereby confirming the analysis of Francis and Ramey (2004). Indeed, also under flexible prices, hours worked would fall significantly as indicated in the upper-right-hand panel of Figure 7. Given these estimates, it is unlikely that a more accommodative monetary policy would lead to positive employment effects. The relatively low medium-run positive effects on hours worked are due to two factors. First, although persistent, the productivity shock is temporary. As a result, output already starts returning to baseline when the effects on hours worked start materializing. A different stochastic process for the productivity shock, which implies a gradual introduction of higher total factor productivity, could increase the effect on hours worked.¹⁶ Second, a positive productivity shock reduces the fixed cost per unit of production, and therefore less labor is required for a given output.

D. The “Great Inflation” and the “Great Moderation”: Subsample Estimates

In this section we first compare the estimates for two subsamples in order to investigate the stability of the full-sample estimates, and then examine, using those estimates, why output and inflation volatility have fallen in the most recent period. The first subsample, corresponding to the period 1966:2–1979:2, captures the period of the “Great Inflation” and ends with the appointment of Paul Volcker as chairman of the Federal Reserve Board of Governors. The second subsample, 1984:1–2004:4, captures the

¹⁵ This picture does not change very much when we do not allow for a positive effect of productivity on exogenous spending.

¹⁶ See, for instance, Rotemberg (2003) for arguments favoring a slow appearance of major productivity advances in output growth.

TABLE 5—SUBSAMPLE ESTIMATES

	Structural parameters				Shock processes				
	1966:1–1979:2		1984:1–2004:4		1966:1–1979:2		1984:1–2004:4		
	Mode	SD	Mode	SD	Mode	SD	Mode	SD	
φ	3.61	1.03	6.23	1.12	σ_a	0.58	0.05	0.35	0.02
σ_c	1.39	0.22	1.47	0.13	σ_b	0.22	0.04	0.18	0.02
h	0.63	0.07	0.68	0.04	σ_g	0.54	0.05	0.41	0.03
ξ_w	0.65	0.07	0.74	0.13	σ_l	0.52	0.09	0.39	0.05
σ_l	1.52	0.65	2.30	0.67	σ_r	0.20	0.02	0.12	0.01
ξ_p	0.55	0.08	0.73	0.04	σ_p	0.22	0.03	0.11	0.01
ι_w	0.58	0.13	0.46	0.16	σ_w	0.20	0.02	0.21	0.03
ι_p	0.45	0.18	0.21	0.09	ρ_a	0.97	0.01	0.94	0.02
ψ	0.34	0.13	0.69	0.11	ρ_b	0.39	0.17	0.14	0.08
Φ	1.43	0.09	1.54	0.09	ρ_g	0.91	0.03	0.96	0.01
r_π	1.65	0.19	1.77	0.29	ρ_l	0.60	0.10	0.64	0.07
ρ	0.81	0.03	0.84	0.02	ρ_r	0.22	0.10	0.29	0.10
r_y	0.17	0.03	0.08	0.05	ρ_p	0.51	0.24	0.74	0.13
$r_{\Delta y}$	0.20	0.03	0.16	0.02	ρ_w	0.96	0.02	0.82	0.15
$\bar{\pi}$	0.72	0.11	0.67	0.10	μ_p	0.46	0.20	0.59	0.18
$\beta^{-1} - 1$	0.14	0.06	0.12	0.05	μ_w	0.84	0.07	0.62	0.17
\bar{l}	0.03	0.62	-0.55	1.21	ρ_{ga}	0.58	0.11	0.39	0.11
$\bar{\gamma}$	0.33	0.04	0.44	0.02					
α	0.19	0.02	0.21	0.02					

Note: SD stands for standard deviation.

more recent period of the “Great Moderation,” in which not only was inflation relatively low and stable, but also output and inflation volatility fell considerably (e.g., Margaret M. McConnell and Gabriel Perez-Quiros 2000). Table 5 compares the mode of the posterior distribution of the DSGE model parameters over both periods.

The most significant differences between the two subperiods concern the variances of the stochastic processes. In particular, the standard errors of the productivity, monetary policy, and price mark-up shocks (and to a lesser extent the investment shock) seem to have fallen. The persistence of those processes has changed much less. One exception is the risk premium shock, which has become even less persistent in the second subperiod.

Somewhat surprisingly, the steady-state inflation rate is only marginally lower in the second subperiod (2.6) versus the first period (2.9). What is different is the central bank’s reaction coefficient to the output gap, which is halved and is no longer significant in the second period. In contrast, the response to inflation is only marginally higher in the second period, and the response to the change in the output gap is the

same. These results are consistent with the findings of Athanasios Orphanides (2003), who shows, using real-time data estimates, that what has changed in US monetary policy behavior since the early 1980s is the relative response to output. They are, however, at odds with the results of Jean Boivin and Marc Giannoni (2006), which find that a stronger central bank response to inflation in the second subperiod can account for a smaller output response to monetary policy shocks estimated in identified VARs. In our case, the lower response to the output gap actually increases the output response of a monetary policy shock in the second period.

Interestingly, it turns out that the degree of price and wage stickiness has increased in the second period, while the degree of indexation has fallen. The latter is consistent with single-equation subsample estimates of a hybrid New Keynesian Phillips curve by Galí and Gertler (1999). This finding is also consistent with the story that low and stable inflation may reduce the cost of not adjusting prices and therefore lengthen the average price duration leading to a flatter Phillips curve. At the same time, it may also reduce rule-of-thumb behavior and indexation

TABLE 6—ACTUAL, MODEL-BASED, AND COUNTERFACTUAL STANDARD DEVIATIONS OF GDP GROWTH AND INFLATION

	1966:1–2004:4		1966:1–1979:2		1984:1–2004:4		Counterfactual 1984:1–2004:4		
	Actual	Model	Actual	Model	Actual	Model	Shocks	Policy	Structure
Growth	0.86	0.94	1.01	1.13	0.59	0.73	1.21	0.70	0.75
Inflation	0.62	0.57	0.55	0.81	0.25	0.34	1.30	0.39	0.32

Notes: “Actual” refers to the data-based standard deviations over the indicated sample; “model” refers to the standard deviations generated by the DSGE model estimated over the indicated sample. The counterfactual standard deviations for the period 1984:1–2004:4 refer to the standard deviations that would have occurred in this period if the shock processes (“shocks”), the monetary policy rule (“policy”), or the structural parameters (“structure”) would have been the same as the ones estimated in the 1966:1–1979:2 sample.

leading to a lower coefficient on lagged inflation in the Phillips curve. The effects are most visible in the goods market, less so in the labor market. Finally, there is also some limited evidence of increased real rigidities in the second subsample. For example, the elasticity of adjusting capital increases from 3.6 to 6.4 in the second subsample.

In order to assess the sources behind the great moderation of the last two decades, Table 6 provides the results of a counterfactual exercise in which we examine what the standard deviation of output growth and inflation would have been in the most recent period if the US economy had faced the same shocks as in the 1970s, if the monetary policy reaction function as estimated in the pre-1979 period would have been the same, or if the structure of the economy would have remained unchanged. Table 6 first of all confirms that both output growth and inflation were significantly less volatile in the second subsample. The estimated DSGE model captures this reduction in volatility, although it overestimates the standard deviation somewhat in both periods. Turning to the counterfactual exercise, it turns out that the most important drivers behind the reduction in volatility are the shocks, which appear to have been more benign in the last period. A reversal to the monetary policy reaction function of the 1970s would have contributed to somewhat higher inflation volatility and lower output growth volatility, but these effects are very small compared to the overall reduction in volatility. Finally, the changes in the structural parameters do not appear to have contributed to a major change in the volatility of the economy. Overall, these results appear to confirm recent findings of James H. Stock and Mark W. Watson (2003)

and Sims and Zha (2006) that most of the structural change can be assigned to changes in the volatility of the shocks. It remains an interesting research question whether policy has contributed to the reduction of those shocks.

VI. Concluding Remarks

In this paper, we have shown that modern micro-founded NNS models are able to fit the main US macro data very well, if one allows for a sufficiently rich stochastic structure and set of frictions. Our results support the earlier approaches by Rotemberg and Woodford (1997) and Christiano, Eichenbaum, and Evans (2005). Although the estimated structural model is highly restricted, it is able to compete with standard VAR and BVAR models in out-of-sample forecasting, indicating that the theory embedded in the structural model is helpful in improving the forecasts of the main US macro variables, in particular at business cycle frequencies.

Of course, the estimated model remains stylized and should be further developed. In particular, a deeper understanding of the various nominal and real frictions that have been introduced would increase the confidence in using this type of model for welfare analysis. Our analysis also raises questions about the deeper determinants of the various “structural” shocks, such as productivity and wage mark-up shocks that are identified as being important driving factors of output and inflation developments. We hope to have shown, however, that the Bayesian approach followed in this paper offers an effective tool for comparing and selecting between such alternative microfounded model specifications.

REFERENCES

- Altig, David E., Lawrence J. Christiano, Martin Eichenbaum, and Jesper Linde.** 2004. "Firm-Specific Capital, Nominal Rigidities, and the Business Cycle." Federal Reserve Bank of Cleveland Working Paper 0416.
- Basu, Susanto, and Miles S. Kimball.** 2002. "Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption." Unpublished.
- Bauwens, Luc, Michel Lubrano, and Jean-François Richard.** 2000. *Bayesian Inference in Dynamic Econometric Models*. Oxford: Oxford University Press.
- Bernanke, Ben S., Mark Gertler, and Simon Gilchrist.** 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." In *Handbook of Macroeconomics Volume 1C*, ed. John B. Taylor and Michael Woodford, 1341–93. Amsterdam: Elsevier Science, North-Holland.
- Bils, Mark, and Peter J. Klenow.** 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy*, 112(5): 947–85.
- Bovin, Jean, and Marc Giannoni.** 2006. "Has Monetary Policy Become More Effective?" Centre for Economic Policy Research Discussion Paper 5463.
- Brooks, Stephen P., and Andrew Gelman.** 1998. "General Methods for Monitoring Convergence of Iterative Simulations." *Journal of Computational and Graphical Statistics*, 7(4): 434–55.
- Calvo, Guillermo A.** 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics*, 12(3): 383–98.
- Chang, Yongsung, João F. Gomes, and Frank Schorfheide.** 2002. "Learning-by-Doing as a Propagation Mechanism." *American Economic Review*, 92(5): 1498–1520.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans.** 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" In *Handbook of Macroeconomics Volume 1A*, ed. John B. Taylor and Michael Woodford, 65–148. Amsterdam: Elsevier Science, North-Holland.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans.** 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113(1): 1–45.
- Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfusson.** 2004. "What Happens after a Technology Shock?" Unpublished.
- Christiano, Lawrence J., Roberto Motto, and Massimo Rostagno.** 2003. "The Great Depression and the Friedman-Schwartz Hypothesis." *Journal of Money, Credit, and Banking*, 35(6): 1119–97.
- Clarida, Richard, Jordi Galí, and Mark Gertler.** 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature*, 37(4): 1661–1707.
- Dedola, Luca, and Stefano Neri.** 2004. "What Does a Technology Shock Do? A VAR Analysis with Model-Based Sign Restriction." Centre for Economic Policy Research Working Paper 4537.
- Eichenbaum, Martin, and Jonas Fisher.** Forthcoming. "Estimating the Frequency of Reoptimisation in Calvo-style Models." *Journal of Monetary Economics*.
- Francis, Neville, and Valerie A. Ramey.** 2005. "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited." *Journal of Monetary Economics*, 52(8): 1379–99.
- Galí, Jordi.** 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89(1): 249–71.
- Galí, Jordi, and Mark Gertler.** 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics*, 44(2): 195–222.
- Galí, Jordi, and Pau Rabanal.** 2005. "Technology Shocks and Aggregate Fluctuations: How Well Does the Real Business Cycle Model Fit Postwar U.S. Data?" In *NBER Macroeconomics Annual 2004*, Volume 19, ed. Mark Gertler and Kenneth Rogoff, 225–88. Cambridge, MA: MIT Press.
- Geweke, John.** 1998. "Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication." Unpublished.
- Goodfriend, Marvin, and Robert G. King.** 1997. "The New Neoclassical Synthesis and the Role of Monetary Policy." In *NBER Macroeconomics Annual 1997*, ed. Ben S. Bernanke and Julio J. Rotemberg, 231–83. Cambridge, MA: MIT Press.
- Kimball, Miles S.** 1995. "The Quantitative Analytics of the Basic Neomonetarist Model."

- Journal of Money, Credit, and Banking*, 27(4): 1241–77.
- King, Robert G., and Sergio T. Rebelo.** 1999. “Resuscitating Real Business Cycles.” In *Handbook of Macroeconomics Volume 1B*, ed. John B. Taylor and Michael Woodford, 927–1007. Amsterdam: Elsevier Science, North-Holland.
- Litterman, Robert B.** 1984. “Forecasting and Policy Analysis with Bayesian Vector Autoregression Models.” *Federal Reserve Bank of Minneapolis Quarterly Review*, 8(4): 30–41.
- McConnell, Margaret M., and Gabriel Perez-Quiros.** 2000. “Output Fluctuations in the United States: What Has Changed since the Early 1980’s?” *American Economic Review*, 90(5): 1464–76.
- Orphanides, Athanasios.** 2003. “Historical Monetary Policy Analysis and the Taylor Rule.” *Journal of Monetary Economics*, 50(5): 983–1022.
- Peersman, Gert, and Roland Straub.** 2004. “Technology Shocks and Robust Sign Restrictions in a Euro Area SVAR.” European Central Bank Working Paper 373.
- Rotemberg, Julio J.** 2003. “Stochastic Technical Progress, Smooth Trends, and Nearly Distinct Business Cycles.” *American Economic Review*, 93(5): 1543–59.
- Rotemberg, Julio J., and Michael Woodford.** 1997. “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy.” In *NBER Macroeconomics Annual 1997*, ed. Ben. S. Bernanke and Julio. J. Rotemberg, 297–346. Cambridge, MA: MIT Press.
- Shapiro, Matthew D., and Mark Watson.** 1988. “Sources of Business Cycle Fluctuations.” In *NBER Macroeconomics Annual 1988*, ed. Stanley Fischer, 111–48. Cambridge, MA: MIT Press.
- Sims, Christopher A.** 2003. “Probability Models for Monetary Policy Decisions.” Unpublished.
- Sims, Christopher A., and Tao Zha.** 1998. “Bayesian Methods for Dynamic Multivariate Models.” *International Economic Review*, 39(4): 949–68.
- Sims, Christopher A., and Tao Zha.** 2006. “Were There Regime Switches in U.S. Monetary Policy?” *American Economic Review*, 96(1): 54–81.
- Smets, Frank, and Raf Wouters.** 2003. “An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area.” *Journal of the European Economic Association*, 1(5): 1123–75.
- Smets, Frank, and Raf Wouters.** 2005. “Comparing Shocks and Frictions in U.S. and Euro Area Business Cycles: A Bayesian DSGE Approach.” *Journal of Applied Econometrics*, 20(2): 161–83.
- Stock, James H., and Mark W. Watson.** 2003. “Has the Business Cycle Changed?” Paper presented at the Monetary Policy and Uncertainty Symposium, Federal Reserve Bank of Kansas City, Jackson Hole, WY.
- Taylor, John B.** 1993. “Discretion versus Policy Rules in Practice.” *Carnegie-Rochester Conference Series on Public Policy*, 39(0): 195–214.
- Woodford, Michael.** 2003. *Interest and Prices*. Princeton: Princeton University Press.

Copyright of *American Economic Review* is the property of American Economic Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.